

REMARKS

This is a response to the Office Action mailed February 22, 2006. Claims 25-30 are pending in this application.

Claim Rejections under 35 U.S.C. 103

In paragraph 4 of the Office Action, claims 25-30 were rejected under 35 U.S.C. 103(a) as obvious over U.S. Patent No. 5,542,073 to Schiefer et al. ("Schiefer"). Applicant respectfully traverses the rejections.

The present application relates to the underlying structure of a database and, in particular, to storing data in new ways that provide an advantage in space usage and/or speed of access over conventional record-based tables. As described in the specification, the underlying data structures include, for example, "instance" and "connectivity" information, where instance information identifies instances of each value in a field that is in a record and connectivity information associates each instance with a specific instance of a value in another field. In certain described embodiments, structures containing cardinality information are used to associate instances with values and/or vice versa.

Schiefer in contrast does not relate to or describe the underlying structures in which data in a database are stored. Instead, Schiefer describes a query optimizer and, in particular, a method for estimating join result sizes used in query optimization. (*See, e.g.*, Schiefer, col. 5, l. 63 - col. 6, l. 62). The query optimization techniques in Schiefer do not require a database having any particular underlying structure and apply to conventionally structured databases. Schiefer simply does not describe or suggest the structures and techniques claimed in the present application.

With respect to claim 25, Schiefer does not disclose or suggest all the limitations of the claim and, accordingly, does not render claim 25 obvious. Claim 25 recites, *inter alia*, a system comprising "a collection of a number of instances corresponding to a value of a first attribute" and "a cardinality element corresponding to the number of instances" of that value of the first attribute. The cardinality element recited in claim 25 thus contains information regarding the number of instances of an attribute having a particular value. Schiefer simply does not disclose such a cardinality element at all. Schiefer instead describes two kinds of statistics, (1) the number of tuples contained in a relation, which Schiefer refers to as "the relation's cardinality," and (2) the number of distinct values taken by an attribute. (Schiefer, col. 1, ll. 39-44). However, neither of these statistics is a cardinality element containing

information regarding the number of instances of an attribute having a particular value or suggests such a cardinality element.

Claim 25 further recites that the cardinality element is “updated each time the number of instances changes.” This limitation recites a linkage between instance information and cardinality information; i.e., each time the number of instances changes, the corresponding cardinality element changes. This linkage was described in the specification in connection with, for example, embodiments of the invention, where changes in instance information result in changes in cardinality elements; the cardinality elements being used in these embodiments to associate instances with values and/or values with instances.

Applicant respectfully disagrees with the statement in Paragraph 4 of the Office Action that “Schiefer teaches that effective cardinality should be determined when the value of particular attribute changes [and that] therefore, it would have been obvious to a person of ordinary skill in the computer² art at the time the invention was [sic] to update cardinality in order to efficiently evaluate the cost estimate to obtain the lowest cost.” Schiefer only suggests determining certain statistics at the time query optimization is performed; it does not describe or suggest updating a cardinality element each time the number of instances changes, as claimed. Applicant also notes that the portion of Schiefer cited by the Examiner, column 3, lines 21-32, contains no discussion of cardinality information or updating cardinality information and respectfully submits that it does not provide a basis for the claim rejection.

Claim 25 also recites that “at least one” of the instances corresponding to a value of a first attribute “indicates at least one other instance corresponding to a value of a second attribute.” In other words, claim 25 requires that at least one instance of a first attribute comprise connectivity information that indicates a corresponding instance of a second attribute. Schiefer does not disclose or suggest such connectivity information; it is simply not directed to such underlying data structures in which data in a database are stored.

Should the Examiner maintain this rejection, Applicant respectfully requests that he specifically identify the portions of Schiefer that disclose each element of the claim.

With respect to claim 26, Schiefer again does not disclose or suggest all the limitations of the claim and accordingly does not render it obvious. Claim 26, like claim 25, recites, *inter alia*, a cardinality element “corresponding to the number of instances” of a value of the first attribute. Schiefer does not disclose this limitation for the same reasons presented with respect to the “cardinality element” of Claim 25. Claim 26 further recites that “the value

[corresponding to an instance] can be derived from the cardinality element.” Schiefer again simply does not disclose this limitation.

Claim 26, again like claim 25, also recites “a collection of a number of instances corresponding to a value of a first attribute” in which “at least one instance indicates at least one other instance corresponding to a value of a second attribute.” Schiefer once again does not disclose this element.

Also, again, the portions of Schiefer cited by the Examiner do not appear to Applicant to be related to this limitation. In particular, Schiefer at column 2, lines 41-56 and column 8, lines 11-19 discusses the effect of “local predicates” on “join result sizes”, which does not relate to the limitations of the claim.

Again, should the Examiner maintain this rejection, Applicant respectfully requests that he specifically identify the portions of Schiefer that disclose each element of the claim.

With respect to claim 27, Schiefer once again does not disclose or suggest all the limitations of the claim and accordingly does not render it obvious. Claim 27 is directed to an aspect of the present invention, described, e.g., on page 22, lines 29-36 of the application, that achieves space savings in the storage of instance elements for two or more tuples when at least two tuples have identical values for at least first and second attributes. In this case, instead of having one instance element for the first attribute of each of the at least two tuples, the system, as recited in claim 27, has a single instance element for the first attribute of at least two such tuples. As further recited in claim 27, this instance element identifies the first attribute value and the second attribute value, and a cardinality element, which as recited in claim 28 may comprise the instance element, identifies a number of tuples having the identical first and second attribute values. Schiefer simply does not disclose such an instance element. Moreover, instance elements, as described in the present application, are part of the underlying data structures used to store the data in a collection of tuples (e.g., a database) and Schiefer is not concerned with and does not disclose such data structures. In addition, Schiefer does not disclose or suggest the claimed cardinality element for all the reasons presented with respect to claims 25 and 26.

With respect to claim 28, it is dependent on claim 27 and is thus patentable for at least the reasons given above in connection with claim 27.

Again, should the Examiner maintain the rejections of claims 27 and 28, Applicant respectfully requests that he specifically identify the portions of Schiefer that disclose each element of the claim.

The Office Action failed to show how Schiefer is applied to claims 29-30 to render them obvious. Nonetheless, applicant respectfully traverses the rejection. Claim 29 recites “a cardinality store storing information representing frequencies of occurrence of instances of equal value, wherein a particular value in the value store associated with a particular instance in the instance store is derived using the cardinality store.” In other words, the cardinality store provides two types of information. First, it provides count-type information “representing frequencies of occurrence of instances of equal value,” that is the number of tuples in the database in which a selected attribute has a particular value. Second, it also provides tuple-identifying-type information for associating “a particular value in the value store . . . with a particular instance in the instance store.” In other words, the tuple-identifying information associates instances in the instance store with values in the “value store.” But since the instances in the instance store represent occurrences of values present in the database tuples, the tuple-identifying information actually relates tuples to their values (hence its name).

Schiefer simply does not disclose such a cardinality store at all. As mentioned above, Schiefer describes only two kinds of cardinality statistics: (1) the total number of tuples contained in a relation, which Schiefer refers to as “the relation’s cardinality;” and (2) the number of distinct values taken by an attribute. (See Schiefer at, e.g., col. 1, lines 39-56). Neither of these statistics is count-type information regarding the number of instances of an attribute having a particular value, or suggests such count-type information. The number of distinct values of an attribute is not the number of instances of particular attribute values. Not only does Schiefer’s cardinality information include the completely different count-type information, but it includes no tuple-identifying-type information at all, either directly or indirectly. Thus, Schiefer does not disclose or suggest the claimed “cardinality store” as recited in claim 29.

Schiefer also does not disclose or suggest an “instance store” storing instances of values that are in turn stored in a separate but associated “value store.” All Schiefer describes are prior art databases having fixed-length tuples with fields where attribute values are stored. (See Schiefer at, e.g., col. 1, lines 23-36). In such databases, the tuple fields storing attribute values are the same as instances; value information is not stored separately from instance information as claimed; and no values are derived.


In sum, claim 29 is patentable over Schiefer and so is its dependent claim 30.

Conclusions

In light of the above remarks, applicant respectfully requests that the Examiner reconsider this application with a view towards allowance. The Examiner is invited to call the undersigned attorney if a telephone call could help resolve any remaining items.

Date: July 21, 2006

Respectfully submitted,



Ognjan V. Shentov

38,051

(Reg. No.)

JONES DAY

222 East 41st Street

New York, New York 10017

(212) 326-3939